

The Power of an Example: Hidden Set Size Approximation Using Group Queries and Conditional Sampling

Dana Ron*

Gilad Tsur†

April 23, 2014

Abstract

We study a basic problem of approximating the size of an unknown set S in a known universe U . We consider two versions of the problem. In both versions the algorithm can specify subsets $T \subseteq U$. In the first version, which we refer to as the *group query* or *subset query* version, the algorithm is told whether $T \cap S$ is non-empty. In the second version, which we refer to as the *subset sampling* version, if $T \cap S$ is non-empty, then the algorithm receives a uniformly selected element from $T \cap S$. We study the difference between these two versions under different conditions on the subsets that the algorithm may query/sample, and in both the case that the algorithm is adaptive and the case where it is non-adaptive. In particular we focus on a natural family of allowed subsets, which correspond to intervals, as well as variants of this family.

*danaron@tau.ac.il, Tel Aviv University. Supported by ISF grant number 671/13.

†gilad.tsur@gmail.com, Tel Aviv University. Supported by the Check Point Institute for Information Security .

1 Introduction

Consider the following problem: For a known universe of elements U and an unknown (“hidden”) subset of it $S \subseteq U$, we wish to approximate the size of S . Of course, our ability to do so is influenced by what access we have to information about S . We consider two basic versions of this problem, where our algorithm has two different ways to obtain information about S :

- **Subset Queries** (also known as **Group Queries**): In this version we specify a subset $T \subseteq U$ and are told whether $S \cap T$ is empty or contains at least one element.
- **Subset Samples**: In this, second, version we again specify a subset $T \subseteq U$, but are given a uniformly selected element of $S \cap T$ if such an element exists, and an indication that the intersection is empty otherwise.

Clearly, subset samples are at least as powerful as subset queries. Note that subset queries (and hence subset samples) can be viewed as a generalization of membership queries (checking whether a single element in U belongs to S). Also note that subset samples can be viewed as a generalization of sampling an element uniformly from S . We study the number of subset queries/samples required to approximate the size of S , where we may be restricted in which subsets $T \subseteq U$ we are allowed to query (or sample). This restriction is on the samples we may perform and not on the structure of S .

Stockmeyer [10, 11], in work better known for results on approximation algorithms for $\#\mathcal{P}$, considers the problem of set size approximation using what we call subset queries.¹ He considers the problem for different families of subsets that may be queried and where the set S may also be restricted to belong to a certain family. We further discuss Stockmeyer’s results (in particular in relation to our results) in Subsections 1.2 and 1.4. We mention briefly that subset queries arise in the context of *Group Testing*, where they are referred to as *group queries* (lending the paper part of its title). Estimation of the size of an unknown set appears in many different settings, including estimating the support size of a distribution and estimating the coverage of search engines. We discuss these settings as well in Subsection 1.4.

1.1 Precise Problem Definition

We say that an algorithm is a *set size approximation algorithm* if, for any set $S \subseteq U$, given an approximation parameter $\epsilon \in (0, 1]$ and either access to subset queries or to subset samples, the algorithm returns an estimate \hat{w} such that with probability at least $2/3$, $\frac{1}{1+\epsilon}|S| \leq \hat{w} \leq (1+\epsilon)|S|$. The success probability of $2/3$ can be increased to $1 - \delta$ for any $\delta > 0$ by standard techniques at a multiplicative cost of $\log(1/\delta)$ in the complexity of the algorithms. For an algorithm performing subset queries, the failure probability is over the coin-tosses of the algorithm, and for an algorithm that uses subset samples, this probability is also over the samples that the algorithm observes. We are interested in the number of queries/samples used by the algorithm as a function of the size of the universe U , which we denote by n , the size of set S , which we denote by w , and the approximation parameter ϵ . This number may be a random variable, in which case we shall bound the probability that it exceeds $g(n, w, \epsilon)$ for a function g that we shall specify. We consider both

¹We cite both the conference version [10] and the journal version [11] of Stockmeyer’s paper, since some of the results appear only in the conference version. Subset queries are given different names in the two versions, including *intersection samples* and *subset samples*.

adaptive algorithms and *non-adaptive* algorithms. In this context, a non-adaptive algorithm cannot determine the next subset it queries/samples based on previous answers to queries/samples, but it can decide when to stop based on the information it obtains. This is necessary in order to allow a dependence on the unknown value w (rather than only on n and ϵ).

1.2 Our Results

We give upper and lower bounds for both versions of our problem and for both non-adaptive and adaptive algorithms in several settings, as detailed next. In what follows we focus on the dependence on n and $w = |S|$. In all cases the dependence on $1/\epsilon$ in the upper bounds is polynomial, and the lower bound holds for constant ϵ (the exact dependencies appear in the statement of each theorem). The results are (generally) asymptotic and we use the $\tilde{O}(f)$ notation to hide a dependency on $\log(f)$. Thus, e.g., $\tilde{O}(\log\log(n))$ may hide a $\log\log\log(n)$ factor.

The simplest case we consider is the most basic and restricted setting in which no strict subsets of U are allowed, namely, the only subset allowed is the whole universe U . Obviously, with subset queries it is not possible to approximate the set size (beyond distinguishing between an empty and non-empty set S). On the other hand, subset samples (which give us uniformly selected elements from S) can provide us with a good approximation when using a number of samples that grows like the square root of w (and this dependence on w can easily be seen to be necessary). As there is only one subset we can access, adaptivity plays no role in this setting.

In the other extreme we may allow the queries/samples to be on *all subsets*. Stockmeyer [10] studied this problem for adaptive subset queries, and gave upper and lower bounds for this case.² We refine these results to be in terms of n, w and ϵ , and not only in terms of n . As we are interested in the difference between the adaptive and non-adaptive setting we also give a non-adaptive version of these results. In the non-adaptive version of this problem the number of queries has a logarithmic dependence on w , whereas in the adaptive case the dependence is doubly logarithmic. These dependencies are tight in the case of subset queries.

Between the two extremes we consider a natural family of subsets, which we refer to as *Interval subsets*. In this case the universe U is fully ordered, and the allowed subsets correspond to intervals of the elements. That is, subsets are of the form $T = \{u_i, u_{i+1}, \dots, u_j\}$ where $U = \{u_1, u_2, \dots, u_n\}$. Interval queries have been studied in the context of group testing: Such a query may, e.g., be to see if a certain part of an electric circuit is faulty, or may be to a subset of ordered test-tubes that is more easily accessed by a robotic arm [7]. All of our positive results for interval subsets extend to the more general case of grids in higher (constant) dimension d where the allowed subsets correspond to d -dimensional sub-grids,³ and to the Boolean hypercube, where the allowed subsets correspond to sub-cubes.

Non-adaptive interval queries can be used to approximate the size of a set in a straightforward manner by using the fact that single elements in the universe are also intervals. This implies an upper bound of $O(n/w)$ on the (non-adaptive) query complexity, and we provide a simple matching lower bound. Hence, non-adaptive interval queries do not offer a real advantage over just being able to query single elements of the universe selected uniformly at random. This situation

²Stockmeyer actually goes further and shows that a significantly smaller set of queries suffices to get similar results. He considers subsets of the elements in U that are defined by a hash function on their index. This allows the family of subsets to be polynomial and not exponential in n .

³In fact, the original motivation for this work arose from our interest in estimating the number of 1-pixels in parts of a two-dimensional image.

changes when adaptivity is allowed. Specifically, one of the results in [11] contains (implicitly) an upper bound for adaptive interval queries of $O(\min(w \log(n), n/w))$. We give a lower bound of $\Omega(\min\{w \cdot \log(n/w^2), n/w\})$, which shows that this is (almost) tight.⁴

Turning to interval samples (in contrast to queries) we get from the results mentioned above that $O(\min\{n/w, \sqrt{w}\})$ such non-adaptive subset samples suffice to approximate the set size. Here too we give a matching lower bound. The power of samples when compared to queries really comes to fore when we present an adaptive algorithm that uses $O(\text{poly}(\log(w)))$ interval subset samples. This algorithm can be adapted to additional settings such as d -dimensional grids and the Boolean hypercube, as we explain in Section 2.

Subsets allowed	Subset Queries		Subset Sampling	
Only U		UB: ∞ LB: ∞		UB: $O(\sqrt{w})$ LB: $\Omega(\sqrt{w})$
Interval ^(*)	NA:	UB: $O(n/w)$ [11] LB: $\Omega(n/w)$	NA:	UB: $[O(\min\{n/w, \sqrt{w}\})]$ LB: $\Omega(\min\{n/w, \sqrt{w}\})$
	A:	UB: $O(w \cdot \log(n))$ [11] LB: $\Omega(\min\{w \cdot \log(n/w^2), n/w\})$	A:	UB: $O(\text{poly}(\log(w)))$ LB:
All	NA:	UB: $\tilde{O}(\log(w))$ LB: $\tilde{\Omega}(\log(w))$	NA:	UB: $[\tilde{O}(\log(w))]$ LB:
	A:	UB: $\tilde{O}(\log \log(w))$ [10] LB: $\tilde{\Omega}(\log \log(w))$ [10]	A:	UB: $[\tilde{O}(\log \log(w))]$ LB:

Table 1: A table of results. ‘NA’ stands for Nonadaptive and ‘A’ for Adaptive. ‘UB’ stands for Upper-bound and ‘LB’ for Lower-bound. The dependence on $1/\epsilon$ in all upper bounds is polynomial. We put brackets around those results that follows directly from other cells in the table, and do not explicitly indicate the $O(n/w)$ upper bound that holds in almost all cases. The results by Stockmeyer are originally as a function of n . (*) The upper bounds for the case of intervals apply to d -dimensional grids where d is a constant and to the Boolean hypercube.

The results described above are summarized in Figure 1. As the results in each row of the table are for refinements of the row above it, one can always use upper bounds from cells above and to the left and lower bounds from cells below and to the right. Obviously lower bounds for adaptive algorithms apply to non-adaptive ones and non-adaptive upper bounds are true also for adaptive algorithms.

1.3 Discussion and Techniques

When reviewing the results described in the previous subsection, the following relationship can be gleaned: As our ability to specify subsets becomes less restricted, the advantage of subset samples over subset queries (i.e., of getting a uniformly selected member of the subset) decreases. An interesting problem is to formalize and prove such a relationship. We conjecture that subset samples do not give any significant speedup when the subsets we have access to are unrestricted. Another general phenomenon (which is of course not unique to the problems we consider), is the

⁴Stockmeyer [11] proves a lower bound for the related problem he studies, but since the allowed query sets in his case are more restricted, the implications on interval queries are not clear (furthermore, he proves his lower bound for sets of a particular size).

power of adaptivity, which comes into play both when using subset queries and when using subset samples.

Whereas some of the algorithms and lower bounds we present are fairly easy to establish, they help to outline the relationship between subset queries and subset sampling. In what follows we highlight two results. The first is the adaptive algorithm that uses interval subset sampling, and which can be adapted to other families of subsets. The second is the lower bound for adaptive interval subset queries. We believe that there are structural aspects of the proof of this lower bound that are interesting.

In the adaptive interval sampling algorithm we maintain a sequence of intervals I_0, \dots, I_t , starting from an initial interval $I_0 = U$, and ending with an interval containing a single element of S (with high probability). The algorithm works iteratively, where in each iteration it continues with a subinterval I_j of the previous interval I_{j-1} . Each new subinterval I_j is selected (using sampling) so that the size of I_j 's intersection with S is a constant fraction (not far from $1/2$) from the size of the intersection of I_{j-1} with S . This ratio between the sizes of the two intersections is then estimated, and the final output is based on the sequence of these estimations.

In the lower bound for adaptive interval queries, we “force” any algorithm to deal separately with roughly w different parts of the universe and “make it” use $\Omega(\log(n/w^2))$ queries for any such part. To formalize this we prove a type of direct-sum claim for a basic problem of determining the location of a single ‘1’ in a binary string. This allows us to come close to matching the upper bound.

1.4 Related Research

As mentioned above, subset queries (in contrast to subset samples) were studied by Stockmeyer [10, 11] in a different context, motivated by complexity theoretical problems (e.g., approximating the number of inputs that satisfy a circuit). Where in the current work we restrict the subsets that the algorithm may query but allow the set S to be arbitrary, Stockmeyer restricts the set itself in addition to restricting the queries. He studies this problem for adaptive algorithms only, and gives results for three families of subsets. The one that is analogous to a case in our work is when neither the subsets that are queried, nor the set S , are restricted (as mentioned above). A second family is where the queries and the set must conform to a certain rooted tree configuration. As mentioned previously, the algorithm given for this case translates directly to our adaptive interval query condition. The third family of subsets is related to matchings in a graph, and does not appear to have direct relevance to our work. All the upper and lower bounds given by Stockmeyer, in contrast to those in this paper, are in terms of n and not of w and ϵ .

Subset queries arise also in the context of *Group Testing*. In group testing, the algorithm is given access to subset queries (referred to as *group queries*) but the goal is to return the set S itself rather than an estimate of the size of S . Research in the field of group testing began in the early 1940's [6] and is still of interest today (e.g. [8]). Many different settings are considered in group testing, with a strong focus on non-adaptive vs. adaptive tests. Similarly to what is done in this work, there is work on group testing when there are limitations on the subsets that the algorithm can query (for more details, see, e.g., [7]). The process we describe can be thought of as a relaxation of group testing, or may be used to get a quick estimate of the number of members in the hidden set, as a preliminary stage for Group Testing.

In the practical setting of estimating the relative coverage of search engines, Bharat et al. [2] consider the problem of estimating the relative size of two sets by performing uniform queries on a

set and then checking whether the elements sampled appear in the other set (and vice-versa). These operations, in our setting, are modeled by performing subset samples on the entire universe and by performing subset samples on single elements. Broder et al. [3] tackle similar problems, and mention the great utility uniform sampling may have for solving them. Anagnostopoulos et al. [1] consider some additional problems related to sampling from search engines in an attempt to estimate sizes. None of the results correspond directly to ours.

Approximating the size of a set using subset samples can be viewed as a special case of approximating the support size of a distribution using conditional sampling. When given access to samples generated according to a general distribution over a domain of size n where each element in the support has probability at least $1/n$, there are almost linear lower bounds for approximating the support size [9, 13, 12]. The best lower bound, due to Valiant and Valiant [12], is $\Omega(n/\log(n))$, and this bound is tight [12]. In the conditional sampling model [5, 4] it is possible to specify a subset T of the domain of the distribution and obtain samples according to the corresponding conditional distribution. The goal is to test various properties of the distribution or estimate various measures. Thus, approximating the size of a set using subset samples corresponds to approximating the support size using conditional sampling in the special case where the distribution is uniform over a subset S .

1.5 Further Research

We suggest several questions and directions for further research.

1. We showed that in the case of interval subsets, adaptive subset sampling significantly improves the complexity of set size approximation: The dependence on w is reduced from linear to polylogarithmic and the (logarithmic) dependence on n is removed. The question is whether the dependence on w can be further reduced (recall that when all subsets are allowed then there is an adaptive algorithm whose complexity is $\tilde{O}(\log\log(w))$).
2. As already mentioned, the upper bounds on the complexity of the problem when using subset sampling with all subsets are based on subset queries. An interesting question is whether sampling has any more power in this setting. Indeed, obtaining lower bounds for subset sampling, in particular for adaptive algorithms, seems more challenging than for subset queries. A related question, which was mentioned earlier as well, is whether it is possible to formalize and analyze the relationship between the generality of the allowed subsets and the power of subset sampling as compared to subset queries.
3. In this work we considered the case of interval subsets and related families of subsets. We believe that there are other families of subsets that may arise naturally in different contexts and are worth studying.

2 Interval Subset Queries and Subset Samples

Since the main focus of this work is on Interval Subsets (and related families of subsets), we start with our algorithms and lower bounds for this family of subsets.

As explained in the introduction, in the case of interval subsets, the domain is a fully ordered set $U = \{u_1, u_2, \dots, u_n\}$, and the subsets that can be queried/sampled, correspond to intervals of the form $\{u_i, u_{i+1}, \dots, u_j\}$. As we discuss at the end of this section, we can extend our results to

d -dimensional grids for constant d , where queries correspond to d -dimensional sub-grids, and to the Boolean hypercube, where queries correspond to sub-cubes.

ON THE STRUCTURE OF OUR LOWER BOUNDS. Before stating precisely and proving our results for interval subsets, we shortly introduce an approach that we take in several of our lower bound proofs. We first assume towards a contradiction that there exists an algorithm that approximates the set size with constant success probability using a certain number of queries $q \leq f(n, w)$ (where n is the universe size and w is the size of the hidden set S). We then define a distribution over pairs of sets S_1, S_2 such that $|S_1| < c|S_2|$ for some constant $c > 0$. Clearly, for every choice of S_1 and S_2 respecting the above conditions, the algorithm we assume exists must be able to distinguish between the hidden set S_1 and the hidden set S_2 with constant probability bounded away from $1/2$ (where the probability is over the coin tosses of the algorithm and in the case of subsets sampling also over the choice of the sampled points). However, from an averaging argument this means that there exists a deterministic algorithm that can distinguish with a probability bounded away from $1/2$ between pairs selected from the aforementioned distribution (where the probability is over the choice of the pair (S_1, S_2) and in the case of subsets sampling also over the choice of the sampled points). Whenever the algorithm performs a subset query or asks for a subset sample from one of the sets, we give it “for free” the result of the query/subset sample on the other set as well. To obtain the lower bound we show that if the algorithm performs fewer than $f(n, w)$ queries (subsets samples), then the answers it gets are equally/similarly distributed, implying that the algorithm cannot distinguish between the two sets with constant probability.

2.1 Interval Subset Queries

Theorem 2.1 *The following holds for the query complexity of set size approximation algorithms that use only interval subset queries:*

1. *There exists a non-adaptive set size approximation algorithm that performs $O(n/(w\epsilon^2))$ interval queries with high constant probability. Furthermore, the probability that the number of queries it performs is larger by a factor of k than this upper bound decreases exponentially with k .*
2. *Any non-adaptive set size approximation algorithm that performs only interval queries performs $\Omega(n/w)$ such queries with constant probability (for constant ϵ). This lower bound holds even when the algorithm is given a constant factor approximation, \tilde{w} of w .*
3. *There exists an adaptive set size approximation algorithm that always performs $O(w \log(n))$ interval queries and with high constant probability performs $O(n/(w\epsilon^2))$ interval queries. Furthermore, the probability that the number of queries it performs is larger by a factor of k than the latter upper bound decreases exponentially with k .*
4. *Any adaptive set size approximation algorithm performs $\Omega(\min(w \log(n/w^2), n/w))$ interval queries with constant probability (for constant ϵ). This lower bound holds even when the algorithm is given a constant factor approximation, \tilde{w} of w .*

Proof: Verifying all the items in Theorem 2.1 aside of Item 4 is fairly simple. We begin by explaining them briefly and then turn to the proof of Item 4.

ITEM 1. The algorithm referred to in this item uses the fact that single elements are a special case of intervals. The algorithm works in two stages. In the first stage it obtains a rough estimate \tilde{w} of $w = |S|$. This is done by iteratively selecting single elements uniformly at random from U and performing subset queries on them. This stage ends once the algorithm receives a positive answer in some iteration j , and \tilde{w} is set to n/j . For any $c > 1$, the probability that $j < n/(c \cdot w)$, so that $\tilde{w} > c \cdot w$ is upper bounded by $(w/n) \cdot n/(c \cdot w) = 1/c$, and the probability that $j > c \cdot (n/w)$, so that $\tilde{w} < w/c$ is upper bounded by $(1 - w/n)^{c \cdot (n/w)} < e^{-c}$. In the second stage of the algorithm it selects $s = \Theta(n/(\tilde{w}\epsilon^2))$ elements uniformly at random, makes a subset query on each, and sets its output, \hat{w} to be (n/s) times the fraction of queries that were answered positively. By the multiplicative Chernoff bound, conditioned on $\tilde{w} \leq 6 \cdot w$ (so that $n/\tilde{w} \geq n/(6w)$), the estimate \hat{w} , is as required with probability at least $5/6$. Since the probability that $\tilde{w} \leq 6 \cdot w$ is at most $1/6$, we get that \hat{w} is as required with probability at least $2/3$. The furthermore claim in this item directly follows from the aforementioned bound on the probability that \tilde{w} underestimates w .

ITEM 2. For the lower bound in this item consider partitioning the universe into intervals of size \tilde{w} , denoted $I_1, \dots, I_{n/\tilde{w}}$. We select a value j^* uniformly at random from $\{1, \dots, n/\tilde{w} - 1\}$. The smaller set S_1 consists of all the elements in I_{j^*} . The larger set S_2 consists of all the elements in $I_{j^*} \cup I_{j^*+1}$. Thus $|S_1| = \tilde{w}$ and $|S_2| = 2\tilde{w}$. Consider any sequence of interval queries T_1, \dots, T_q where $q < n/(12 \cdot \tilde{w})$. The probability (over the choice of j^*), that for some set T_i , at least one of its endpoints belongs to I_{j^*} or I_{j^*+1} is at most $1/3$. But if such an event does not occur, then the answers to the queries T_1, \dots, T_q are the same for S_1 and S_2 and they cannot be distinguished.

ITEM 3. The algorithm referred to in this item combines two procedures. The first procedure, which performs $O(w \log(n))$ queries (deterministically), determines the set S exactly. This is done by performing a type of “extended” binary search. That is, the search constructs a binary tree, where each node in the tree corresponds to an interval query. The root of the tree corresponds to all of U . If a query on an interval I is answered positively where $|I| > 1$, then the corresponding node has two children, one for the left half of I and one for the right. A node is a leaf if either the answer on the corresponding interval is negative, or if the answer is positive and the interval is of size 1. Thus, we have a non-empty leaf for every element of S , and the total number of queries performed is $O(w \log n)$. The second procedure runs the non-adaptive algorithm of Item 1 to estimate the size of S . By performing interleaved queries and stopping when the first procedure stops the desired bound is obtained.

ITEM 4. For this lower bound we first address the (simpler) case in which $\tilde{w} \geq \sqrt{n}$ (so that the lower bound should be $\Omega(n/w)$). The construction is the same as in Item 2, except that here S_1 and S_2 also include the first element from each interval I_j . Since the number of intervals is $n/\tilde{w} \leq \tilde{w}$, we have that $|S_1|$ and $|S_2|$ are both $\Theta(\tilde{w})$ and $|S_2|$ is a constant factor larger than $|S_1|$. By this construction, any query that does not contain *only* elements of a single interval I_j , is answered positively both for S_1 and for S_2 . On the other hand, as long as the algorithm does not ask a query $T \subset I_j$ for $j \in \{j^*, j^* + 1\}$, it will get a negative answers for both S_1 and S_2 . The lower bound follows.

It remains to address the case that $\tilde{w} < \sqrt{n}$. To gain intuition, consider the following “game”. There are b identically looking locked boxes where only one of these boxes is non-empty. We would like to open the non-empty box, but opening each box requires time t , and until it is open we do not know whether it is empty or not. If the non-empty box is selected uniformly at random, then, with high constant probability, it will take us time $\Omega(bt)$ to get to the non-empty box.

To obtain a lower bound following the above intuition, we partition U into \tilde{w} intervals $I_1, \dots, I_{\tilde{w}}$ of size n/\tilde{w} each (each of these corresponds to a “locked box”). Similarly to the construction for $\tilde{w} \geq \sqrt{n}$, both S_1 and S_2 include the first element from each interval. In addition, we further partition each interval I_j into $m = \lceil n/\tilde{w}^2 \rceil$ subintervals I_j^1, \dots, I_j^m of size \tilde{w} each (more precisely, of size at least $\tilde{w}/2$ and at most \tilde{w}). For both S_1 and S_2 , we select, uniformly at random, one subinterval $I_j^{\ell(j)}$ from each interval j , and both sets include the first and the last element from each of these subintervals. In addition, we select a “special” interval I_{j^*} uniformly at random, and include in S_2 *all* elements in the subinterval $I_{j^*}^{\ell(j^*)}$. Hence, $|S_1|$ and $|S_2|$ are both $\Theta(\tilde{w})$ and $|S_2|$ is a constant factor larger than $|S_1|$.

Consider any (possibly adaptive) deterministic algorithm. We would like to show that in order to distinguish correctly between S_1 and S_2 with probability at least $2/3$ (over the choice of S_1 and S_2), it must perform at least $\tilde{w} \cdot \log(n/\tilde{w}^2)/6$ queries with constant probability. Consider a *process* that answers queries of an algorithm while it selects the pair (S_1, S_2) on the fly (according to the aforementioned distribution). Actually, the process makes all the random choices for S_1 and S_2 , except the choice of j^* . That is, it selects $\ell(j)$ for each j in advance, but does not select j^* .

For each interval query T , if T is not strictly contained within one of the intervals I_j , then the process answers positively for both S_1 and S_2 . For each interval query T that is contained within an interval I_j , if T is not contained within the subinterval $I_j^{\ell(j)}$, then the process gives the same answer for both S_1 and S_2 (i.e., a positive answer if $I_j^{\ell(j)} \cap T \neq \emptyset$ and a negative answer if $I_j^{\ell(j)} \cap T = \emptyset$). Once the algorithm performs a query T that is contained in $I_j^{\ell(j)}$, which we’ll refer to as *revealing* $\ell(j)$, then the process does the following. Let r be the number of indices j for which the algorithm has not yet revealed $\ell(j)$ (where initially $r = \tilde{w}$). With probability $1/r$ the algorithm sets $j^* = j$ (possibly implying a different answer for S_1 and S_2 , so that the algorithm “wins”). Otherwise it increases r by 1 and gives the same answer for both sets (positive if T contains at least one of the end-points of $I_j^{\ell(j)}$ and negative otherwise).

It follows that in order to distinguish between S_1 and S_2 with probability at least $2/3$, the algorithm must reveal at least a one-third of the $\ell(j)$ s with probability at least $1/3$. It remains to show that such a task requires at least $\tilde{w} \cdot \log(n/\tilde{w}^2)/6$ queries with constant probability. In order to obtain this lower bound, we consider a closely related problem, which we refer to as *multiple single elements* that we define next. Each instance of this problem is parameterized by two integer parameters, b and m . It consists of b binary strings of length m each, where in each string s^j , there is a single 1 in a position $\ell(j)$, and all other bits are 0. An algorithm may ask, for any substring t that is a (consecutive) substring of some s^j whether it is all 0 or contains a 1. The goal of the algorithm is to determine (*reveal*) $\ell(j)$ for at least $b/3$ of the substrings s^j . We next prove that any algorithm for the *multiple single elements* problem, when given as input a uniformly selected random instance (i.e., in which each $\ell(j)$ is selected uniformly at random in $\{1, \dots, m\}$), must perform at least $b \cdot \log(m)/c$ queries with high constant probability (for a sufficiently large constant c). This will imply the lower bound claimed in this item, since an algorithm that distinguishes between S_1 and S_2 with probability at least $2/3$ can be used for this problem.⁵

For any algorithm that solves the multiple single elements problem, consider the decision tree that corresponds to it. That is, each internal node in the tree corresponds to a query, and each leaf corresponds to an output $\ell(j_1), \dots, \ell(j_{b/3})$. The number of leaves at depth at most $b \cdot \log(m)/6$ is

⁵The (straightforward) reduction is achieved by mapping between elements of each string s^j and the subintervals I_j^ℓ of I_j .

$m^{b/6}$. For any fixed leaf, the fraction of instances for which it contains a correct answer is $m^{-b/3}$. Therefore, the total fraction of instances for which some leaf at depth at most $b \cdot \log(m)/6$ provides a correct answer is $m^{-b/6}$, and the lower bound follows. ■

2.2 Interval Subset Samples

Theorem 2.2 *The following holds for the sample complexity of set size approximation algorithms that use only interval subset samples:*

1. *Any non-adaptive set size approximation algorithm that performs only interval samples requires $\Omega(\min(n/w, \sqrt{w}))$ such samples.*
2. *There exists a non-adaptive set size approximation algorithm that uses $O(\min(n/w, \sqrt{w})/\epsilon^2)$ interval samples. The probability of using more samples decreases exponentially with the number of samples.*
3. *There exists an adaptive set size approximation algorithm that uses $\tilde{O}(\log(w)^4/\epsilon^2)$ interval samples. The probability of using more samples decreases exponentially with the number of samples.*

We note that the $\log^4(w)$ factor in Item 3 can be reduced to $\log^{3+\gamma}$ for any constant γ , but for the sake of simplicity, we give the slightly higher upper bound.

Proof: The lower bound argument for Item 1 is similar to that used for Item 2 in Theorem 2.1. Consider partitioning the universe into intervals of size \tilde{w} , denoted $I_1, \dots, I_{n/\tilde{w}}$. We select a j^* uniformly at random from $\{1, \dots, n/\tilde{w}\}$. The small set S_1 is composed of $\tilde{w}/2$ elements selected uniformly at random from I_{j^*} . The large set S_2 consists of all the elements in I_{j^*} . Consider any fixed choice of intervals T_1, \dots, T_q for $q < \min(n/\tilde{w}, \sqrt{\tilde{w}})/c$ for a sufficiently large constant c . The probability, over the choice of j^* , that any one of these intervals has at least one end-point in I_{j^*} is at most $1/6$. If this event does not occur, then each interval either contains I_{j^*} or is completely disjoint from it. In the latter case, no sample is returned for both S_1 and S_2 . In the former case, as long as a collision (repetition of an element) does not occur, for both S_1 and S_2 , each new sampled element is uniformly distributed in I_{j^*} . By the upper bound on q , the probability that a collision occurs is a small constant.

The correctness of Item 2 is based on running two (subset query) algorithms in parallel, alternating between them in the choice of intervals. The algorithm that performs $O((n/w) \cdot \epsilon^{-2})$ queries is described in Theorem 2.1, Item 1. The algorithm that performs $O(\sqrt{w} \cdot \epsilon^{-2})$ queries is described in Theorem 3.1. When either algorithm returns a result we return that as our result. The correctness and bounds on the number of samples follow directly from those in the descriptions of the algorithms.

We now turn to describing the algorithm referred to by Item 3. The basic approach in this algorithm is as follows. The algorithm constructs a sequence of intervals I_0, \dots, I_t , where $I_0 = U$, $I_j \subset I_{j-1}$ for each $j \geq 1$ and with high constant probability, $|S \cap I_t| = 1$. Let $b_j \stackrel{\text{def}}{=} |S \cap I_{j-1}|/|S \cap I_j|$ and observe that $|S| = |S \cap I_0| = |S \cap I_t| \cdot \prod_{j=1}^t b_j$. For each pair of intervals I_j and I_{j-1} , the algorithm maintains an estimate \hat{b}_j of b_j such that $1/(1 + \epsilon_j) \leq \hat{b}_j/b_j \leq 1 + \epsilon_j$ for a sufficiently small ϵ_j , and with sufficiently high probability $1 - \delta_j$. The output of the algorithm is $\prod_{j=1}^t \hat{b}_j$. The error parameter ϵ_j and the confidence parameter δ_j should be such that $\prod_{j=1}^t (1 + \epsilon_j) \leq (1 + \epsilon)$ and $\sum_{j=1}^t \delta_j$ is a small constant.

Given a pair of intervals $I_j \subset I_{j-1}$, if $1/4 \leq 1/b_j \leq 3/4$, then, it follows from the multiplicative Chernoff bound that an estimate \hat{b}_j as described above can be obtained by asking $\ln(3/\delta_j)/\epsilon_j^2$ subset samples on the subset I_{j-1} . Therefore, it remains to explain how each I_j is selected based on I_{j-1} so as to ensure (with probability at least $1 - \delta_j$) that indeed $1/b_j$ is as desired. This is done by performing $s_j = 4\ln(3/\delta_j)$ subset samples on the subset I_{j-1} , and ordering the selected elements $v_1 \leq v_2 \leq \dots \leq v_{s_j}$. If they are all equal then the algorithm sets $t = j - 1$. Otherwise, if $I_{j-1} = [w_{j-1}, w'_{j-1}]$, then $I_j = [w_{j-1}, v_{s_j/2}]$. The probability that either the algorithm terminated with I_t such that $|I_t \cap S| > 1$ or that $b_j < 1/4$ or $b_j > 3/4$ is at most δ_j .

If we set $\delta_j = 1/(10j^2)$ and $\epsilon_j = \epsilon/(100j^{3/2})$ then with probability at least $2/3$ the algorithm terminates after $c \log(w)$ iterations (for a constant $c > 1$) with an estimate as required.⁶ In such a case, the total number of subset samples used is $c' \log(w)^4 \log \log(w)/\epsilon^2$ for a constant c' . Since the probability that the algorithm does not terminate in $c \log(w)$ iterations (since it does not obtain an interval of size 1) is at most $1/3$, the probability that it does not terminate in $k \cdot c \log(w)$ iterations is $\exp(-k)$, as desired. ■

2.3 Other Related Families of Subsets

In this subsection we describe how to modify our most efficient algorithm, which uses interval subsets (i.e., the adaptive subset sampling algorithm), to two additional settings, thus obtaining algorithms with complexity $\text{poly}(\log(w))$ in these settings as well. The non-adaptive algorithms using interval subsets work as is for the above two universes and families of subsets (since these algorithms are based on having access to singleton subsets, and to all the universe, which also holds in these settings). The adaptive algorithm using interval queries can be easily modified to construct a “search-tree” whose non-empty leaves contain single elements of the set S .

d -DIMENSIONAL GRIDS AND SUB-GRID SUBSETS. Let U be a hypergrid $\{1, \dots, k\}^d$ where d is a constant, and let the family of subsets that the algorithm can query/sample, correspond to d -dimensional sub-grids. Thus, $n = k^d$, and the special case of $d = 1$ corresponds to interval subsets over a fully ordered universe U . The modified (adaptive subset-sampling) algorithm defines a sequence of sub-grids, R_0, \dots, R_t where $R_0 = U$, $R_{j+1} \subset R_j$ for every $j \geq 0$ and $|R_t \cap S| = 1$ (with high constant probability). The main observation is that for each $j \geq 0$, there exists a sub-grid R_{j+1} such that $|R_j \cap S|/(2d) \leq |R_{j+1} \cap S| \leq (1 - 1/(2d))|R_j \cap S|$, and that a sub-grid with similar properties can be found efficiently (with sufficiently high probability) by sampling. For simplicity, we first establish this observation for $d = 2$, and later explain how it generalizes to larger values of d .

Given a two-dimensional sub-grid R_j whose lower-left corner is (x_j^{\min}, y_j^{\min}) and whose upper-right corner is (x_j^{\max}, y_j^{\max}) , consider all sub-grids defined by “cutting” R_j along the x -axis. That is, sub-grids defined by (x_j^{\min}, y_j^{\min}) and $(x_j^{\min} + b, y_j^{\max})$ or by $(x_j^{\min} + b, y_j^{\min})$ and (x_j^{\max}, y_j^{\max}) (for $1 \leq b \leq x_j^{\max} - x_j^{\min} - 1$). If one of these sub-grids contains at least one-fourth of the points in $R_j \cap S$, then we are done. Otherwise, there must be a value b^* such that the one-dimensional sub-grid defined by $(x_j^{\min} + b^*, y_j^{\min})$ and $(x_j^{\min} + b^*, y_j^{\max})$ contains at least half of the points in $R_j \cap S$. But this implies that there exists a sub-grid defined by (x_j^{\min}, y_j^{\min}) and $(x_j^{\max}, y_j^{\min} + b')$ that contains at least one-fourth of the points in $R_j \cap S$.

⁶The setting of δ_j ensures that the sum over all δ_j , which is the failure probability of the algorithm, converges to a constant. The setting of ϵ_j ensures that the product over all j of $(1 + \epsilon_j)$ is upper bounded by $(1 + \epsilon_j)$.

The algorithm asks for $\Theta(\log(1/\delta_j))$ subset samples with the subset R_j . By the above discussion and a multiplicative Chernoff bound, with probability at least $1 - \delta_j$ there is either a sub-grid defined by (x_j^{\min}, y_j^{\min}) and $(x_j^{\min} + \hat{b}, y_j^{\max})$ that contains between one-eighth and two-eighths of the sample points, or there exists such a sub-grid defined by (x_j^{\min}, y_j^{\min}) and $(x_j^{\max}, y_j^{\min} + \tilde{b})$. The algorithm lets R_{j+1} be this sub-grid. As in the case of interval subsets, the algorithm estimates the ratio between $|R_{j+1} \cap S|$ and $|R_j \cap S|$ and uses the estimates to compute its final output.

The argument generalizes to $d > 2$ by applying an iterative process that considers the d dimensions one after the other, “losing” at most a fraction of $1/d$ of the points in each iteration. In order to ensure that with high probability the algorithm selects a “good” sub-grid in each iteration, the setting of δ_j should be reduced by a factor of d , and the size of the sample required to estimate the ratio between $|R_{j+1} \cap S|$ and $|R_j \cap S|$ needs to be increased by another factor of d (as we need to take a union bound over the d dimensions). Since d is assumed to be a constant, the complexity remains as in the case of interval subsets.

THE BOOLEAN HYPERCUBE AND SUB-CUBE SUBSETS. Let $U = \{0, 1\}^d$ for $d = \log n$ and let the family of allowed subsets consist of all sub-cubes (i.e., subsets of U that are determined by restricting a subset of the coordinates to a fixed value in $\{0, 1\}$). The first basic observation here is that given a sub-cube C such that $|C \cap S| \geq 2$, it contains a sub-cube C^* such that $|C \cap S|/3 \leq |C^* \cap S| \leq 2|C \cap S|/3$. To verify this, assume, without loss of generality, that the restricted coordinates of C are $\{1, \dots, j\}$. Starting from $j + 1$, we restrict the unrestricted coordinates, where we always select the restricted value for which the intersection with S is larger. We stop once we obtain a sub-cube C^* as specified (where we must reach such a stopping condition based on the restriction procedure). In addition, let C^+ be the minimal sub-cube that satisfies $C^* \subset C^+ \subseteq C$ and such that $|C^+ \cap S| \geq 5|C \cap S|/6$.

The second observation is the following. Suppose we ask for $\Theta(\log(1/\delta))$ subset samples with the subset C , and consider the *maximal* sub-cubes C' of C that are defined by restricting coordinates $1, \dots, t$ for some $t \geq 1$ and that contain between $1/4$ and $3/4$ of the sample points. With probability at least $1 - \delta$, one of these sub-cubes will contain C^* and be strictly contained in C^+ . Furthermore, by the maximality of these sub-cubes and the condition on the number of sample points that they contain, there are at most two such sub-cubes. If there is just one, then we are done. Otherwise, we ask for $\Theta(\log(1/\delta))$ additional subset samples with the subset C , and among the two sub-cubes, select one that contains between $1/4$ and $7/8$ of the sample points. In this manner we can obtain (with high constant success probability) a sequence of sub-cubes $C_0 \supset C_1 \supset \dots \supset C_t$ where $C_0 = U$, $|C_t \cap S| = 1$, and $|C_{j+1} \cap S|$ is a constant fraction of $|C_j \cap S|$. By estimating these fractions, we obtain an estimate of S , similarly to the case of intervals.

3 Querying and Sampling From U with no Access to Subsets

It is clear that subset queries (in contrast to subset sampling) provide us almost no information in this setting. We can perform a single subset query (on all of U), and if it returns 1 we know that $|S| \geq 1$. Otherwise, we know that $|S| = 0$. Approximating $|S|$ beyond this is impossible. It remains to discuss sampling. Both the upper and the lower bound are based on the probability of obtaining the same element twice (a *collision*) as a function of the size of the set, w .

Theorem 3.1 *1. There exists a set size approximation algorithm that is provided (only) with samples from U and with high constant probability uses $O(\sqrt{w}/\epsilon^2)$ samples. Furthermore, the*

probability that the number of queries it performs is larger by a factor of k than this upper bound decreases exponentially with k .

2. Every set size approximation algorithm that is provided only with samples from U must use $\Omega(\sqrt{w})$ samples with constant probability for any constant ϵ . This lower bound holds even when the algorithm is given an estimate \tilde{w} such that $\tilde{w}/4 \leq w \leq 4\tilde{w}$.

Proof: We begin with the lower bound. For simplicity we give the argument for $\epsilon = 1$. This directly implies the lower bound for any $\epsilon < 1$. For any given $4 \leq \tilde{w} \leq |U|/4$, consider the following two distributions over sets S . The first distribution is uniform over subsets of size $\tilde{w}/4$ and the second distribution is uniform over subsets of size $4\tilde{w}$. For both distributions, the probability that a sample of size $s = \sqrt{\tilde{w}/c}$ contains some repeated element (a collision) is at most $\binom{s}{2} \cdot \frac{4}{\tilde{w}}$ (the term $\binom{s}{2}$ is the number of sample pairs among s samples, and the term $4/\tilde{w}$ is the probability that a specific pair is a collision in the first distribution (and an upper bound on this probability in the second distribution)). For $c \geq 6$ this probability is at most $1/3$. Conditioned on there being no collisions, the samples in both cases are identically distributed.

Turning to the upper bound, the algorithm works in two stages. In the first stage it obtains an estimate \tilde{w} such that $\tilde{w}/c \leq w \leq c \cdot \tilde{w}$ with probability at least $5/6$ for some constant $c > 1$ (e.g., $c = 100$ suffices). In the second stage it takes a sample of size $\Theta(\sqrt{\tilde{w}}/\epsilon^2)$ to obtain a $(1 + \epsilon)$ -factor approximation. We note that while describing the algorithm as having two stages may seem to imply that it is adaptive, this is not the case. Since the algorithm only receives samples from U , the only decision it makes is when to stop and output an estimate. Details follow.

To obtain the (rough) estimate \tilde{w} , the algorithm takes samples from U until a collision occurs. If this first collision occurs on the j^{th} sample, then the estimate \tilde{w} is set to j^2 . The probability that $j < \sqrt{w/c}$ is upper bounded by $\binom{\sqrt{w/c}}{2} \cdot \frac{1}{w} < 1/c$ (using a union bound on the collision probability of each pair of elements). On the other hand, the probability that $j > \sqrt{c \cdot w}$ is upper bounded by $(1 - \sqrt{c \cdot w}/(2w))^{\sqrt{c \cdot w}/2} < \epsilon^{-c/4}$. This upper bound follows by considering a partition of the sample into two parts of equal size $\sqrt{c \cdot w}/2$. The probability that no collision occurs in the sample is upper bounded by the probability that no collision occurs in the first half of the sample, which is upper bounded by 1, times the probability that no collision occurs between an element selected in the second half and an element selected in the first half (conditioned on the elements in the first half being distinct).

It follows that with probability at least $5/6$ \tilde{w} is within a factor of 6 from w . Furthermore, the probability that the total number of samples taken in the first stage exceeds $k \cdot \sqrt{w}$ decreases exponentially with k . Assume from this point on that indeed $w/6 \leq \tilde{w} \leq 6w$. Also assume that $\epsilon \leq 1/2$, or else set ϵ to $1/2$.

In the second stage the algorithm takes a sample of size $s = \Theta(\sqrt{\tilde{w}}/\epsilon^2)$. For each $1 \leq i < j \leq s$, let $\eta_{i,j}$ be a random variable indicating whether the j^{th} sample is the same as the i^{th} . Thus $\Pr[\eta_{i,j} = 1] = 1/w$, and $\text{Exp}[\sum \eta_{i,j}] = \binom{s}{2} \cdot \frac{1}{w}$. Let $\eta = \sum \eta_{i,j}$. We set \hat{w} to be $\binom{s}{2}/\eta$. By Chebyshev's inequality,

$$\Pr\left[\left|\eta - \text{Exp}[\eta]\right| > (\epsilon/2) \cdot \text{Exp}[\eta]\right] < \frac{4\text{Var}[\eta]}{\epsilon^2 \text{Exp}[\eta]^2}. \quad (1)$$

Now, $\text{Var}[\eta] = \text{Exp}[\eta^2] - \text{Exp}[\eta]^2$ where

$$\begin{aligned} \text{Exp}[\eta^2] &= \sum_{i < j} \text{Exp}[\eta_{i,j}^2] + 4 \sum_{i < j < k} \text{Exp}[\eta_{i,j}\eta_{j,k}] + 6 \sum_{i < j < k < \ell} \text{Exp}[\eta_{i,j}\eta_{k,\ell}] \\ &= \binom{s}{2} \cdot \frac{1}{w} + 4 \binom{s}{3} \cdot \frac{1}{w^2} + 6 \binom{s}{4} \cdot \frac{1}{w^2}. \end{aligned} \quad (2)$$

Since $\text{Exp}[\eta]^2 = \left(\frac{s}{2}\right)^2 \cdot \frac{1}{w^2}$, which is lower bounded by $6 \binom{s}{4} \cdot \frac{1}{w^2}$, we get that $\text{Var}[\eta] = O(s^2/w + s^3/w^2)$. It follows that $4\text{Var}[\eta]/(\epsilon^2 \text{Exp}[\eta]^2) = O(\epsilon^{-2}(w/s^2 + 1/s))$. By the choice of s this is at most $1/6$ (for an appropriate constant in the $\Theta(\cdot)$ notion for s). Therefore, with probability at least $2/3$ we get that $\hat{w} \leq (1 + \epsilon/2)w$ and $\hat{w} \geq (1 - \epsilon/2)w \geq w/(1 + \epsilon)$, as required. As for the sample size, it exceeds $k \cdot \sqrt{w}/\epsilon^2$ only due to \tilde{w} overestimating w . The probability of this event (as described above) decreases exponentially with k . ■

4 Unrestricted Subset Queries and Subset Samples

In this section we prove the next two theorems, and get the corollary that follows. As mentioned in the introduction, analogues of these results with regard to the adaptive case were proved by Stockmeyer [11], who did not relate the complexity to w or ϵ but rather only to n . We provide all details for the sake of consistency.

- Theorem 4.1** *1. There exists a non-adaptive set size approximation algorithm that performs $\tilde{O}(\log(w)/\epsilon^3)$ (unrestricted) subset queries with high constant probability. Moreover, for any integer k , the probability that the algorithm performs a number of queries that is more than a factor of k larger than the above upper bound decreases exponentially with k .*
- 2. There exists an adaptive set size approximation algorithm that performs $\tilde{O}(\log\log(w)/\epsilon^3)$ queries with high constant probability. Moreover, for any integer k , the probability that the algorithm performs a number of queries that is more than a factor of k larger than the above upper bound decreases exponentially with k .*

Stockmeyer further proved that a much smaller family of subsets, polynomial in n of size, may be used to achieve similar results. As a directly corollary of Theorem 4.1 we get that the same upper bounds hold when the algorithm may perform subset samples. We comment that for constant ϵ , Item 1 in Theorem 4.1 is implied by Item 2, since any adaptive subset-query algorithm can be emulated by a non-adaptive subset-query algorithm at an exponential cost in the query complexity. However, since we are interested in a polynomial dependence on $1/\epsilon$, we address the two cases separately.

Theorem 4.2 *Every non-adaptive set size approximation algorithm must perform $\tilde{\Omega}(\log(w))$ queries with probability at least $1/6$ (for any constant ϵ).*

Corollary 4.3 *Every adaptive set size approximation algorithm must perform $\tilde{\Omega}(\log\log(w))$ queries with probability at least $1/6$ (for constant ϵ).*

In the other settings we study, the lower bounds we presented held even if the algorithm was given a constant factor approximation \tilde{w} of the size of S . In contrast, when we are allowed to query arbitrary subsets of U , given such an estimate \tilde{w} , it is possible to obtain a $(1 + \epsilon)$ -approximation by performing a number of queries that depends only on $1/\epsilon$. Hence, the issue is essentially to find such an estimate \tilde{w} , and this is where the dependence on w comes into play. The upper bounds we suggest are based on starting with a small estimated value $\tilde{w} = 1$ and increasing the estimate. In a similar manner one could start with $\tilde{w} = n$ and iteratively decrease it, leading to an upper bound of $\tilde{O}(\log(n/w)/\epsilon^3)$. We could take the minimum between these two values of upper bounds, but then one may start estimating w from, e.g., \sqrt{n} . Indeed for every relation between w and n there is an algorithm that performs $q(\epsilon)$ queries for this particular relation. Therefore, in the current setting, a lower bound of $\Omega(g(w))$ means that every algorithm must perform $\Omega(g(w))$ queries for most values of w .

Proof of Theorem 4.1: Both the non-adaptive and the adaptive algorithm search for an estimate for w by working in iterations, where in each iteration they hold a hypothesis for such an estimate and they test it. The reason for the exponential difference between the complexities is that in the non-adaptive case we perform a “doubling search” on the size of the set, while in the adaptive case we perform both a doubling search and a binary search on \log the size of the set. We first describe the algorithms for the special case of $\epsilon = 1$, and later explain how to modify them to other values of ϵ .

THE NON-ADAPTIVE ALGORITHM (FOR $\epsilon = 1$). For any integer i let $e_i \stackrel{\text{def}}{=} 2^{i-1}$, and let $\delta_i \stackrel{\text{def}}{=} 1/(10i^2)$. Also define $\rho_i \stackrel{\text{def}}{=} (1 - 1/e_i)^{e_i}$ (so that for $i \geq 2$, $1/4 \leq \rho_i \leq 1/e$). The algorithm works in iterations. In iteration i , for each $1 \leq j \leq t_i = c \log(1/\delta_i)$ (where c is a constant) it selects a subset T_i^j by including each element from U in T_i^j independently with probability $1/e_i$ and it performs a subset query on T_i^j . Let \hat{p}_i denote the fraction of subsets T_i^j that return a negative answer (among the t_i subsets queried). The estimate output by the algorithm is e_i for the first i such that $\hat{p}_i \in [\rho_i^{\sqrt{2}} - 0.02, \rho_i^2 - 0.02]$.

For each i let $p_i(w)$ denote the probability (over the choice of each S_i^j) that the query on T_i^j is answered negatively. By the process of selecting each T_i^j we have that $p_i(w) = (1 - 1/e_i)^w = \rho_i^{w/e_i}$. If $e_i < w/2$, then $p_i(w) < \rho_i^2$, and if $e_i > 2w$, then $p_i(w) > \rho_i^{1/2}$. On the other hand, there exists an index $i(w)$ such that $w/\sqrt{2} \leq e_{i(w)} \leq \sqrt{2}w$, implying that $\rho_{i(w)}^{\sqrt{2}} < p_{i(w)}(w) \leq \rho_{i(w)}^{1/\sqrt{2}}$. By the setting of the t_i s (for a sufficiently large constant c), for each (fixed choice of) i , the probability that $|\hat{p}_i - p_i(w)| > 0.02$ is at most δ_i . Summing over all i , the probability that such an event occurs for some i is upper bounded by $1/3$. Since $\rho_i^{\sqrt{2}} - \rho_i^2 > 0.04$ and $\rho_i^{1/2} - \rho_i^{1/\sqrt{2}} > 0.04$ for every i , we get that the probability that the procedure outputs e_i such that either $e_i > 2w$ or $e_i < w/2$ is at most $1/3$, as required. In each iteration i the algorithm performs $O(\log(1/\delta_i)) = O(\log(i))$ queries, and by the above analysis, the probability that it performs more than $k \cdot c \log(w) \log \log(w)$ queries (for a fixed constant c and any k) decreases exponentially with k .

THE ADAPTIVE ALGORITHM (FOR $\epsilon = 1$). The adaptive algorithm works in two stages. In the first stage, rather than increasing its “hypothesis” for w (i.e., e_i) by a factor of 2 in each iteration, the adaptive algorithm increases the \log of the estimate (i.e., i) by a factor of 2 in each iteration. Namely, for each integer $\ell \geq 0$, in iteration ℓ it tests the hypothesis that $e_{2^\ell} \geq w$ and stops in the first iteration ℓ in which the test passes. Similarly to the non-adaptive case, this is done by

selecting $t_\ell = c \log(1/\delta_\ell)$ subsets S_ℓ^j (for $\delta_\ell = 1/(10\ell^2)$), where each element is included in S_k^j with probability $1/e_{2^\ell}$ and performing a subset query on each subset. The fraction of subsets on which a negative answer is returned is denoted \hat{p}_ℓ . The test is said to pass if $\hat{p}_\ell > \rho_{2^\ell}^{1/2} - 0.01$. Similarly to the analysis of the non-adaptive case, with probability at least $5/6$, the test stops in iteration ℓ^* such that $w \leq e_{2^{\ell^*}} \leq w^2$. That is, $\log(w) \leq 2^{\ell^*} \leq 2\log(w)$, so that $\log\log(w) \leq \ell^* \leq \log\log(w) + 1$.

In its second stage, the algorithm performs a binary search for $i(w)$ (where $i(w)$ is as defined in the analysis of the non-adaptive algorithm) between $i_{\min} = 2^{\ell^*-1}$ and $i_{\max} = 2^{\ell^*}$. Each step of the binary search computes an estimate \hat{p}_i as in the non-adaptive algorithm until obtaining an index i such that $\hat{p}_i \in [\rho_i^{\sqrt{2}} - 0.02, \rho_i^2 - 0.02]$. The analysis of the quality of the estimate is as in the non-adaptive case.

DEALING WITH $\epsilon < 1$. In the non-adaptive algorithm and in the second stage of the adaptive algorithm we set $e_i \stackrel{\text{def}}{=} \lfloor (1 + \epsilon/4)^{i-1} \rfloor$ and $t_i = c \log(1/\delta_i)/\epsilon^2$ (the first stage of the adaptive algorithm remains unchanged). Since for any constant x and $\gamma < 1$ we have that $x^{1+\gamma} - x = \Omega(\gamma)$, we modify the rule for the selected index i (in the non-adaptive algorithm and in the second stage of the adaptive algorithm) to be $\hat{p}_i \in [\rho_i - \epsilon/c', \rho_i^{1+\epsilon/4} + \epsilon/c']$ (for an appropriate constant $c' > 1$). The analysis is adapted in a straightforward manner, and adds an additional factor of $O(1/\epsilon)$ due to the modified definition of e_i . ■

As mentioned previously, this upper bound can be “optimized” for different values of w in relation to n . If we, e.g., assume that $w > n/w$, then we can search for w starting from the hypothesis that $w = n$, in an order where each hypothesis is smaller than the previous one. To make this concrete, consider the non-adaptive algorithm for $\epsilon = 1$. In each iteration i it selects sets R_i^j , where each element in U is included in R_i^j with probability $1/e_{\log(n)-i}$. Let \hat{q}_i be the fraction of sets R_i^j on which the subset query returned a negative answers. If $\hat{q}_i \in [\rho_{\log(n)-i}^{\sqrt{2}} - 0.02, \rho_{\log(n)-i}^2 - 0.02]$, then the algorithm outputs $e_{\log(n)-i}$ as its estimate. In this manner, if $w > n/w$, then the search reaches $e_{i(w)}$ after $\log(n/w) < \log(w)$ steps.

We now turn to the lower bounds.

Proof of Theorem 4.2: Here too we take the approach introduced at the beginning of Section 2. Consider the following distribution over pairs of sets (S_1, S_2) . First we select an integer i uniformly at random in $\{0, \dots, \log(n) - 1\}$. The set S_1 is a uniformly selected set of size 2^i and the set S_2 is a uniformly selected set of size 2^{i+1} . We make two simple observations. The first is that if a query is on a set T such that $|T| > 4i \cdot (n/2^i)$, then the probability that the answer to either S_1 or S_2 is negative is $\exp(-\Omega(i))$. The second is that if a query is on a set T such that $|T| < n/(4i \cdot 2^i)$, then the probability that the answer to either S_1 or S_2 is positive is at most $1/(c' \cdot i)$. It follows that if the algorithm performs less than i/c'' queries all of which are on sets of size greater than $4i \cdot (n/2^i)$ or smaller than $n/(4i \cdot 2^i)$, then the probability that the algorithm can distinguish between S_1 and S_2 is a small constant (less than $1/6$).

Consider any fixed sequence of queries T_1, T_2, \dots (where the only decision of the (non-adaptive) algorithm is after which query T_q to stop and give its output). We shall say that a subset T_j is *useful* for i if $j \leq i/(12 \log(i) + 24)$ and $n/(4i \cdot 2^i) \leq |T_j| \leq 4i \cdot (n/2^i)$. The second condition is equivalent to $\log(n) - i - (\log i + 2) \leq \log |T_j| \leq \log(n) - i + (\log i + 2)$. Since $i < \log(n)$ so that $\log(i) < \log\log(n)$, a set T_j can be useful for at most $2\log\log(n) + 4$ values of i . Therefore, the subsets T_1, \dots, T_q for $q \leq \log(n)/(12\log\log(n) + 24)$ can be useful for less than $\log(n)/6$ values of i . Since i is selected uniformly at random in $\{0, \dots, \log(n) - 1\}$, the probability (over the choice of i)

that there is a subset T_j that is useful for i is at most $1/6$. The lower bound on any non-adaptive algorithm follows by combining this with the first part of the proof. ■

References

- [1] A. Anagnostopoulos, A. Z. Broder, and D. Carmel. Sampling search-engine results. *World Wide Web*, 9(4):397–429, 2006.
- [2] K. Bharat and A. Z. Broder. A technique for measuring the relative size and overlap of public web search engines. *Computer Networks*, 30(1-7):379–388, 1998.
- [3] A. Z. Broder, M. Fontoura, V. Josifovski, R. Kumar, R. Motwani, S. U. Nabar, R. Panigrahy, A. Tomkins, and Y. Xu. Estimating corpus size via queries. In *Proceedings of the Fifteenth ACM International Conference on Information and Knowledge Management (CIKM)*, pages 594–603, 2006.
- [4] C. Cannone, D. Ron, and R. Servedio. Testing equivalence between probability distributions using conditional samples. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1174–1192, 2014.
- [5] S. Chakraborty, E. Fischer, Y. Goldhirsh, and A. Matsliah. On the power of conditional samples in distribution testing. In *Proceedings of the Second Symposium on Innovations in Theoretical Computer Science (ITCS)*, pages 561–580, 2013.
- [6] Robert Dorfman. The detection of defective members of large populations. *The Annals of Mathematical Statistics*, 14(4):436–440, 12 1943.
- [7] D. Du and F. Hwang. *Combinatorial group testing and its applications*. World Scientific, 1993.
- [8] P. Indyk, H. Q. Ngo, and A. Rudra. Efficiently decodable non-adaptive group testing. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1126–1142, 2010.
- [9] S. Raskhodnikova, D. Ron, A. Shpilka, and A. Smith. Strong lower bounds for approximating distributions support size and the distinct elements problem. *SIAM Journal on Computing*, 39(3):813–842, 2009.
- [10] L. J. Stockmeyer. The complexity of approximate counting (preliminary version). In *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing (STOC)*, pages 118–126, 1983.
- [11] L. J. Stockmeyer. On approximation algorithms for $\#P$. *SIAM Journal on Computing*, 14(4):849–861, 1985.
- [12] G. Valiant and P. Valiant. Estimating the unseen: an $n/\log(n)$ -sample estimator for entropy and support size, shown optimal via new CLTs. In *Proceedings of the Forty-Third Annual ACM Symposium on the Theory of Computing (STOC)*, pages 685–694, 2011. See also ECCC TR10-179 and TR10-180.

- [13] P. Valiant. Testing symmetric properties of distributions. *SIAM Journal on Computing*, 40(6):1927–1968, 2011.